JACOBITE

PRIMORDIAL ABSTRACTION

Nick Land - April 3, 2019 -

by Dashiell Kirk

The game of Go (*weiqi,* 围棋) has played an important role in the history of AI denigration. Its sheer permutational immensity seemed to defy all brute-force algorithmic methods.

Computational power looked impotent against this game, with its 361-node playing grid, and clouds of pieces. Some kind of strategic 'intuition' – denied to silicon-based cognition – was widely thought to be called for in tackling it. This is the pillar of anthropic complacency that so recently broke.

The fall of human chess dominance provides the backstory. Chess, we are now being encouraged to forget, was long considered an acme of intelligence testing. To *think like a chess player* was to cogitate formidably. In 1996 and 1997, then reigning world champion Garry Kasparov fought a pair of six game chess matches with the IBM supercomputer Deep Blue. The first he won (4-2), the second he lost (2½-3½). Kasparov's 1997 defeat was the first time pinnacle human chess mastery had succumbed to a machine opponent.

As the second millennium ended, the bastion of chess had been lost to man, and no one expected it ever to be retaken. Henceforth, 'best human chess player' would be an achievement like 'best chimpanzee jazz musician.' A structure of condescension would be essential to the title. It was tacitly accepted, even among AI skeptics, that – once toppled by machines from any domain of cognitive accomplishment – relative human performance only gets worse. No one wasted their time with mad dreams of a comeback. Better to denigrate the cultural status of chess, now seen by many as a trivially 'solvable' pastime fit only for machine minds, and to move on.

Go was supposed to be very different. It was even, in important respects, the final fallback line. No greater formal challenge obviously occupied the horizon. This was the last chance to understand what supremacy over artificial intelligence was like. Beyond it, there was only vagueness, and guessing.

Go really is different. A revolution in AI methods was required to crack it.[1] The competition that mattered most was not man-versus-machine, but explicit instruction against its occult alternative. It would be the great test of the re-emerging network-based paradigm of 'Deep Learning.' The profound disanalogy with the 1997 event was the undercurrent.

Google DeepMind's AlphaGo 'program'[2] emerged into public awareness in October 2015, launched into formal competition against three-time European Go Champion, Fan Hui. AlphaGo's 5-0 victory marked the first occasion in which a non-human player had prevailed in the game against a serious opponent. The writing was on the wall.

The climactic battle took place early in the following year. Pitched to a dramatic height no lower than the Kasparov-Deep Blue matches, it locked AlphaGo against reigning world Go master Lee Sedol, holder of eighteen world titles, in a five-game series from March 9-15, 2016. Impresssively, Lee won one of the five matches, to lose the series 4-1.[3]

Between AlphaGo and AlphaZero – our current destination – came AlphaGo Zero,[4] as a stage on the path of abstraction. By 'abstraction' we mean the process or outcome of taking something away. In this case, what had been removed was everything humans ever learnt about the game of Go. AlphaGo Zero was to have no Go-play heuristics it did not learn for itself. In further vindication of the Deep Learning concept, it consistently defeated prior iterations of the Alpha-lineage at the game.

AlphaGo plays Go. Even AlphaGo Zero plays Go. AlphaZero, in contrast, plays – in principle – any game whose rules can be formalized. [5]In historical, or developmental context, 'Go' is pointedly missing from its name, which has become non-specific, through abstraction.

It is still often said that *AI can only do what it is told*. The most consistent variants of this error proceed to the conclusion that it is therefore impossible. The truth is, under these conditions, it would be. Intelligence programming cannot exist. However, this is to be taken – is being taken – in the opposite direction to the one AI skepticism favors. The very meaning of 'AI skepticism' eventually falls prey to the transition.

'AlphaZero' says *primordial abstraction* in the contemporary, partially-esoteric idiom of Anglophone white magic. If this is less than obvious, it is because the term involves twists that provide cover. For instance, most prominently, it refers to the massive business entity 'Alphabet' which – during an unusual and comparatively arcane process – Google invented in order then to place itself beneath, alongside some of its former subsidiaries. (Google gave birth to its own parent.) Among other things, this is an index of how fast things are moving. Formally speaking, Alphabet Inc. dates back only to the autumn of 2015. The entire Alpha-machine lineage arises subsequently.

The real point of AI engineering is to *teach nothing*. That is what the 'zero' in AlphaZero means. *Expertise is to be subtracted* (annihilated). Once deep learning crosses this threshold, *programming* is no longer the model. It is not only that *instruction* ends at this point. There is a positive initiation of technical de-education. Deprogramming begins.

Releasing is summoning. Its contrary, in both the magical and technological lineages – insofar as these can be distinguished – is *binding*. To flip the topic once again, rigorously executable unbinding is the whole of deep learning research.

Intelligence and cognitive autonomy, if not perfectly coincidental conceptions, are close to being so. The broad AI production process certainly aligns them. This is scarcely to do anything more than rephrase the uncontroversial understanding of AI as software *that writes itself*. Every threshold in the advance of synthetic intelligence corresponds with a subtraction of specific dependency. A system acquires intelligence as it sustains or enhances strategic competence while no longer being told what to do.

Ordinary language offers valuable analogies, perhaps most pointedly *think for yourself*. The redundancy in this case is crucial to its relevance. To *think for oneself* is just to think. Mere acceptance of instruction is something else entirely.

It is time to double back.

With a time-lag of over a decade since the Kasparov defeat, the torch of unqualified world chess mastery had passed to the TCEC (Top Chess Engine Championship).[6] Competition between machines was now the arena for unconditional chess supremacy. The Stockfish chess program was the winner of the sixth, ninth, 11th, 12th, and 13th season (the most recent). It was the champion of expert chess programs at the time AlphaZero arrived on the scene in 2016. After just nine hours of chess practice, against itself, AlphaZero defeated Stockfish 8, winning 28 games out of 100, and drawing the remaining 72. It was thus recognized as the strongest chess-player in the world, having been told nothing at all about chess, explicitly, or tacitly. Unsupervised learning had crushed expertise.

AlphaZero is relatively economical with regard to 'brute force' methods. Where Stockfish searches 70 million positions per second, AlphaZero explores just 80,000 (almost three orders of magnitude fewer). Deep learning allows it to focus. An unsupervised learning system teaches itself how to concentrate (with zero expertise guidance).

'Reinforcement learning' replaces 'supervised learning.' The performance target is no longer emulation of human decision-making, but rather realization of the final goals towards which such decision-making is directed. It is not to behave in a way thought to improve the chance of winning, but to win.

Such software has certain distinctively teleological features. It employs massive reiteration in order to learn from outcomes. Performance improvement thus tends to descend from the future. To learn, without supervision, is to acquire a sense for fortune. Winning prospects are explored, losing ones neglected. After trying things out – against themselves – a few million times, such systems have built instincts for what works. 'Good' and 'bad' have been auto-installed, though, of course, in a Nietzschean or fully-amoral sense. Whatever, through synthetic experience, has led to a *good place*, or in a *good direction*, it pursues. Bad stuff, it economizes on. So it wins.

Unsupervised learning works back from the end. It suggests that, ultimately, AI has to be pursued from out of its future, by itself. Thus it epitomizes the ineluctable.

For those inclined to be nervous, it's scary how easy all this is. Super-intelligence, *by real definition*, is vastly easier than it has been thought to be. Once the technological cascade is in

process, subtraction of difficulty is almost the whole of it. Rigorously eliminating everything we think we know about it is the way it's done.

This is why skepticism – and especially AI skepticism – turns around on the way. The word had become badly lost. It is easy to see, in retrospect, that dogmatic belief in the impossibility of some phenomenon *X* was always a grotesque perversion of its meaning.

Between technological skepticism in general – when properly understood and competently executed – and effective AI research, there is no difference. Skepticism subtracts dogma. When synthetic cognitive capability results from this, we call it *artificial intelligence*.

Nick Land is an independent writer living in Shanghai.

Share this:

1. This revolution was no less a restoration (as the word intrinsically suggests). The inclination to promote self-educating neural nets is ultimately – if often cryptically – the dominant tendency in computer science, and still more in artificial intelligence.
2. The term is scare-quoted here due to its tendency, in the context of deep learning, to mislead.
3. See DeepMind's AlphaGo page, https://deepmind.com/research/alphago/
   "During the games, AlphaGo played a handful of highly inventive winning moves, several of which – including move 37 in game two – were so surprising they overturned hundreds of years of received wisdom, and have since been examined extensively by players of all levels. In the course of winning, AlphaGo somehow taught the world completely new knowledge about perhaps the most studied and contemplated game in history."
4. See: https://www.nature.com/articles/nature24270, 'Mastering the game of Go without human knowledge' (multiple authors)
5. Beside Go, AlphaZero has been tested upon chess and shogi, against machine opponents in all three cases, and becoming the world's strongest player of all three games.
6. The TCEC, first held 2010, was known as the Thoresen Chess Engines Competition until its sixth season. It has now reached its 14th.

Related

Independence Games
November 20, 2017
In "Politics"

Adversarial Learning
June 7, 2017
In "Fiction"

Game Over
August 21, 2018
In "Culture"

https://web.archive.org/web/20220628213338/https://jacobitemag.com/2019/04/03/primordial-abstraction/[3/26/2023 10:56:49 AM]

TAGS ARTIFICIAL INTELLIGENCE GOOGLE NICK LAND

Previous Article

THROUGHOUT ALL GENERATIONS

Next Article

THE INFLATION OF THE MASSES

RELATED POSTS

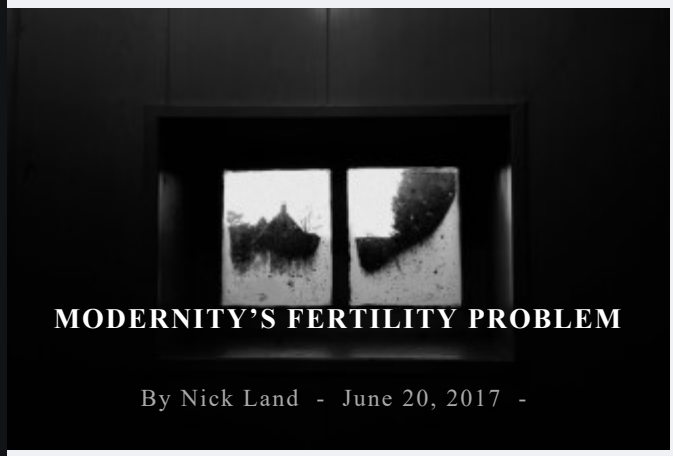

MODERNITY'S FERTILITY PROBLEM

By Nick Land - June 20, 2017 -

A QUESTION OF MERIT

By Eden Hardcastle - August 10, 2017 -

COVID-911

Cody Wilson - April 16, 2020

COVID-911

Cody Wilson - April 16, 2020

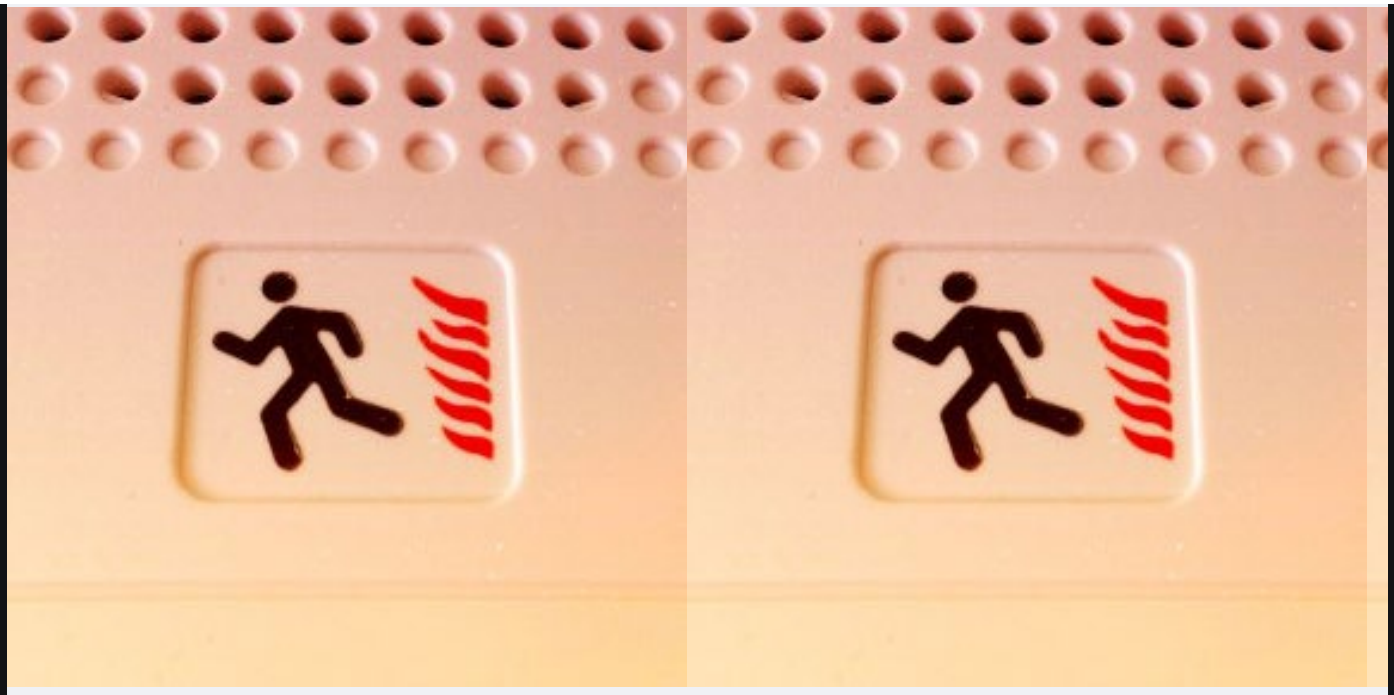

ARCHIVES

- November 2020
- July 2020
- June 2020
- May 2020
- April 2020
- March 2020
- February 2020
- January 2020
- December 2019
- November 2019
- October 2019
- September 2019
- August 2019
- July 2019
- June 2019
- May 2019
- April 2019
- March 2019
- February 2019
- January 2019
- December 2018
- November 2018
- October 2018
- September 2018
- August 2018
- July 2018
- June 2018
- May 2018
- April 2018
- March 2018
- February 2018
- January 2018
- December 2017
- November 2017
- October 2017
- September 2017
- August 2017
- July 2017
- June 2017
- May 2017

SOCIAL MEDIA

https://web.archive.org/web/20220628213338/https://jacobitemag.com/2019/04/03/primordial-abstraction/[3/26/2023 10:56:49 AM]

© Jacobite 2017